2023 27th International Conference Information Visualisation (IV) | 979-8-3503-4161-4/23/$31.00 ©2023 IEEE | DOI: 10.1109/IV60283.2023.00021

Visualizing Maps of Visitors' Interest for Museum Exhibits with Single-Board Computers

Shigeo Takahashi,* Yohei Nishidate,* Yukihide Kohira,* and Rentaro Yoshioka*
*University of Aizu, Aizu-Wakamatsu, Japan

***Abstract*—Understanding the degree of satisfaction for visitors has been a key factor in selecting attractive collections and designing appealing layouts in art galleries and museums. Although monitoring the actual spatiotemporal behaviors of visitors is essential for this purpose, introducing an expensive monitoring system would impose a heavy burden on the financial management and leads to unwanted restrictions on the layout design in the exhibition rooms. This paper presents an approach to visualizing the spatiotemporal changes in the maps of visitors' interest with a system of installed single-board computers such as Raspberry Pi devices. Employing single-board computers as IoT sensors facilitates monitoring systems to maximally covers the entire exhibition space while keeping the associated installation cost and power consumption sufficiently low. Our approach for this novel system organization begins by first detecting individuals from camera images using machine learning techniques and reconstructing their spatial positions from perspective views. Kernel density estimation was employed to represent the distribution of interest across the entire exhibition room as a continuous function by respecting the reconstructed positions of visitors. This allowed the use of heatmaps to visualize the changes in the map of interest reflecting the travel history of individual visitors and the accumulated distribution of interest over a specific period. Experimental results from eight months of measurement data demonstrate the capability of the proposed approach, including meaningful trends that reveal how the layout of collections attracted visitors to the exhibitions.**

***Index Terms*—maps of interest, single-board computers, spatiotemporal changes, heatmaps, exhibition layout design**

I. Introduction

The spatial placement of collections significantly impacts the attractiveness of individual pieces of work, especially in art galleries and museums. Curators often extract historical relations between such pieces before finding their optimal layout to enhance visitors' understanding of the underlying backgrounds of collections in the exhibition. However, they usually have *qualitative* means of assessing the spatial design of these exhibits only, for example, by asking visitors to participate in survey questionnaires to obtain their feedback. This consideration leads us to seek an effective tool for *quantitatively* evaluating the goodness in the spatial placement of pieces in the exhibition space.

The approach explored in this study is to track the behavior of visitors in the exhibition room so that we can understand their interest and preference in the exhibited pieces of work. This is analogous to understanding the preference of individuals for visual media by tracking the movement of their gaze points, for example. Nonetheless, recording the movements of visitors usually involves the use of expensive sensor devices or

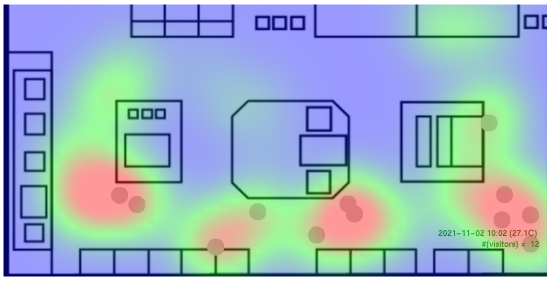

Fig. 1. Visualizing spatiotemporal changes in the map of visitors' interest as a heatmap.

tracking markers equipped with communication devices. This inevitably results in unwanted design constraints on the layout of the exhibition space and imposes physical sensor devices on visitors. It is also crucial to avoid collecting visitors' personal information when tracking their spatiotemporal positions.

In this paper, we present an approach for composing *maps of interest* to represent how visitors are interested in the collection of exhibits. Our tools for this approach are inexpensive single-board computers equipped with sensors, such as *Raspberry Pi* devices, which facilitate tracking the spatiotemporal behavior of visitors in exhibition rooms. The bounding boxes of visitors in the camera images are first extracted using a machine-learning-based object detection technique, and these are subsequently used to reconstruct their 3D positions by inverting the viewing transformation. By synchronizing the camera frames of multiple sensors, we compose an entire layout of visitors in the exhibition space. We evaluate such layouts as heatmaps by visualizing the spatiotemporal distributions of visitors' interests as dynamic density distribution maps. As experimental results, we accumulate the distribution maps of interest frame by frame as an integral map to understand how the overall visitors distributed their interest in individual pieces of work during specific periods. Fig. 1 shows an example in which we visualize spatiotemporal changes in the map of visitors' interest as a heatmap.

This study's contributions can be summarized as follows:

- Reasonably affordable single-board computers are employed as IoT sensors, using their programmable characteristics to process the measurement data.
- A novel approach designed explicitly for single-board computers is developed to compose maps of interest by tracking the spatiotemporal behaviors of museum visitors.

2375-0138/23/$31.00 ©2023 IEEE
DOI 10.1109/IV60283.2023.00021

The rest of this paper is organized as follows. Section II provides a brief survey on related topics in the context of museum exhibition design and visualization. Section III explains the setup of single-board computers as IoT sensors we installed in the museum. Section IV describes how we reproduced the spatiotemporal behavior of visitors in the museum using techniques for machine-leaning-based object detection and 3D reconstruction. Section V details our techniques for visualizing maps of visitors' interests in museum exhibits. Section VI presents our analysis of the viewing behaviors of visitors from eight months of measurement data, followed by a discussion of the proposed approach. Finally, we conclude this paper and refer to possible future extensions in Section VII.

II. Related Work

We briefly review previous work on designing exhibition layouts and visualizing the distributions of interest.

A. Designing Exhibition Layouts

Much research has been conducted to evaluate the selection and layout of exhibits in museum spaces.

Choi [1] explored ways to understand the relationship between the spatial layout of the museum and the exploration patterns of visitors, and Hillier and Tzortzi [2] analyzed the spatial configuration of the exhibition to understand its impact on the formation of traveling visitors. Peponis et al. [3] further studied how the thematic grouping of exhibits influences the spatial behaviors of visitors. Yalowitz and Bronnenkant [4] focused on the timing and tracking of museum visitors to extract their characteristic traveling patterns from observation. Bitgood [5] surveyed pieces of literature on visitor circulation in exploring the museums, and Kirchberg and Tröndle [6] reviewed studies on visitors' experiences in exhibitions and identified similarities and dissimilarities between them. Yoshimura et al. [7] employed Bluetooth sensors to collect anonymized data about visitors to the Louvre Museum. In particular, they focused on the microscopic behaviors of visitors in exploring the museum and unveiled the underlying mechanism of congestion in the physical space. Simulating the exploration of museum exhibits in augmented-reality environments was proposed in [8].

B. Visualizing Distributions of Interest

Visualization techniques help us not only to facilitate the understanding of historical and scientific contents exhibited in the museums [9], [10] but also to interpret the movements of visitors exploring the exhibition space.

Visualizing distributions of interest in the exhibition space relies on the analysis of movement data [11]. Several pioneering studies were done by Andrienko et al. [12], [13], in which they successfully extracted meaningful features from detailed movement data by visualizing it as a set of trajectories. Due to recent advancements in sensing technologies, several indoor positioning approaches have become available for our use [14]. Lanir et al. [15] introduced a proximity-based indoor positioning system equipped with radio frequency identification (RFID) tags [16] to track the movements of individual visitors and sophisticated visualization tools to enhance the visual readability of their exploration behaviors in the museum. Visualization approaches [17], [18] also benefit us in producing continuous spatiotemporal changes by interpolating relatively discrete samples in terms of time. Another tool for interactively navigating moving patterns of characters facilitates an understanding of the behavioral trends of visitors to the exhibition space [19]. Extracting specific patterns from the observed behavior of museum visitors [20] and visitor pairs [21] have also been tackled. Note that understanding human movements through sound [22] may allow us to implement multimodal visualization models.

The aim of this study is to visualize the spatiotemporal behaviors of visitors to encode the dynamic changes in the population density distribution in the exhibition room. For this purpose, we can employ heatmap representations by considering a visitor traveling in the exhibition room as a gaze point moving on the screen. Visualizing the movements of eye gazes has been essential to evaluating the quality of visual information and its associated interfaces [23], and heatmaps along with gaze plots have often been used for this purpose [24], [25]. In this study, we employ a technique for visualizing time-varying heatmaps proposed in [26] to illuminate the underlying trends in the spatiotemporal behaviors of museum visitors.

Fig. 2. Raspberry Pi board equipped with sensors, including a camera.

III. Installing Sensors

This project is carried out in collaboration with the Fukushima Museum, a prefectural museum in Aizu-Wakamatsu city. The museum allows visitors to view exhibits freely if they pay a relatively low admission fee. In this section, we describe how single-board computers were installed as sensors in the exhibition room of the museum.

A. Installing Raspberry Pi Devices as Sensors

Raspberry Pi devices were selected as single-board computers because they are readily available and cost-effective for educational purposes. In practice, the Raspberry Pi computer allows us to accommodate sensors as its peripherals, including a thermometer, a hygrometer, an illuminance meter, a sound level meter, a pressure gauge, a camera, a thermal imaging camera, a motion detector, and so on. In this study, photos captured by the Raspberry Pi 4 with the camera device were used to track the positions of visitors. These Raspberry Pi

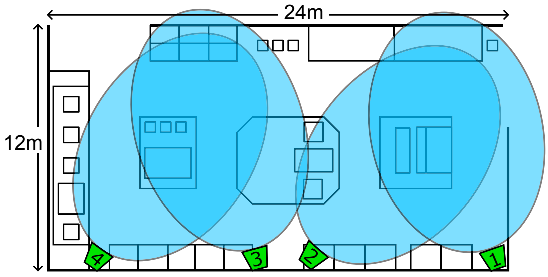

Fig. 3. Sensors installed in the exhibition room.

Fig. 4. Detecting visitors with the YOLO algorithm.

devices send measurement data via the local wireless network available exclusively for this project and save them on the disk storage maintained by the server. Fig. 2 shows the Raspberry Pi device that was installed in the exhibition room.

Several software programs were installed on the Raspberry Pi computers to conduct the necessary pre-processing of measurement data. The device does not send actual photos captured by the camera to avoid storing identity information about the visitors. Instead, the pre-processed results obtained by the installed programs were saved. Each single-board device was programmed to store the measurement data every minute during the opening hours of the museum.

B. Sensor Setup

Four Raspberry Pi sensors were installed in the exhibition room for collections from the Tumulus period. Fig. 3 illustrates the layout of the sensors installed in the room, where we tried to cover the entire space with the four sensors by adjusting their camera directions. The sensors were placed at an approximate height of 2.5 m to avoid disturbing the visitors' views. In the experiments, the measurement data were analyzed offline as a post-process to reproduce the spatiotemporal changes in the map of interest over the exhibition room.

IV. Tracking the Positions of Visitors

This section describes the approach used to retrieve the 3D positions of visitors in the exhibition room.

A. Detecting Visitors in Camera Images

The first task was identifying visitors in the camera images captured by the respective sensors. For this purpose, the *YOLO* algorithm [27] was employed, allowing us to detect objects and their classification types. The YOLO algorithm detects such specific objects in images and videos with trained deep *convolutional neural networks*, and its precision is high compared to other algorithms. Fig. 4 demonstrates an example of persons detected as bounding boxes using the YOLO algorithm.

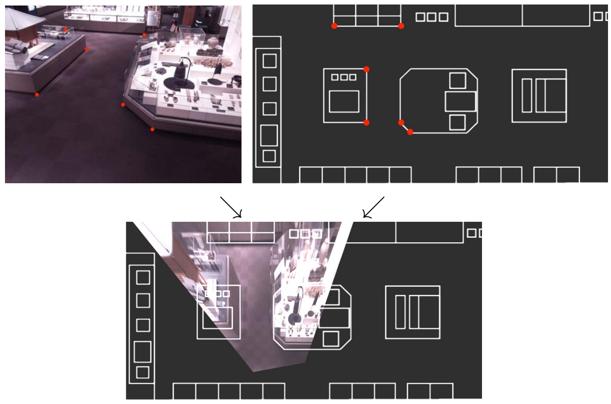

Fig. 5. Computing a homography from a single perspective image to the ground plane by matching pairs of corresponding points.

The *YOLOv3* algorithm [28] for detecting visitors as persons was implemented using Python and installed on the Raspberry Pi computer. This implies that the camera images captured by the Raspberry Pi sensors were never directly stored. Instead, the corner coordinates of the bounding boxes that enclose the detected persons as visitors were recorded. In this way, personal information about visitors identified from the camera images can immediately be discarded.

B. Reconstructing the Positions of Visitors

Having extracted the bounding boxes of visitors in the exhibition room, their standing positions on the ground were explored. The aim was to eliminate perspective distortion in the camera image by calculating a *homography* from the perspective image to the ground plane of the floor plan. For further information, refer to additional details about homography in several technical papers [29], [30].

To carry out this approach, homography computation tools in the OpenCV library were employed, where four or more pairs of corresponding points in the two images needed to be matched [31]. Fig. 5 depicts a case where six pairs of corresponding points were manually plotted between a single perspective image and the floor plan image. This manual plotting of matching points for the single perspective image captured by each sensor was carried out to compute the standing positions of visitors.

It was still necessary to identify the ground position of each visitor in the perspective view from the corresponding bounding box obtained by the YOLO algorithm. For each bounding box, the midpoint of the bottom edge of the box may be employed as the standing point of the visitor. However, it is sometimes the case that some exhibits occlude the lower body of a visitor, and thus the bounding box is truncated to enclose the upper body only. This problem was resolved by

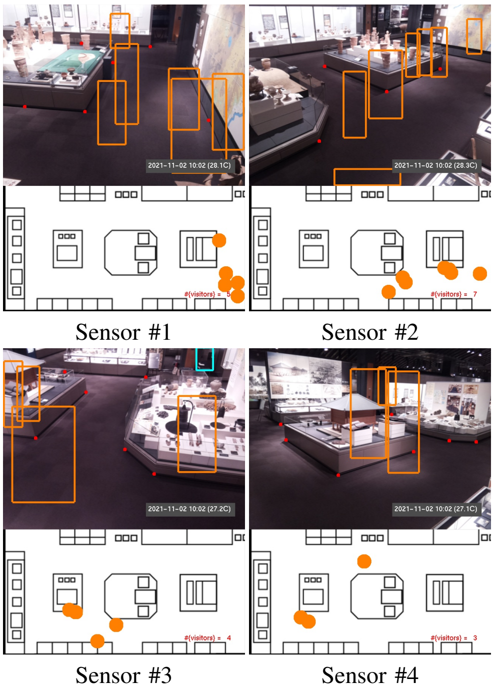

Fig. 6. Detected bounding boxes (in orange) and manually plotted feature points (in red) in the camera view of each sensor and reconstructed positions of the visitors (as orange disks) in the exhibition room.

computing the aspect ratio of each bounding box. Empirically, we observed that the ratio of the box width to the height is around 1:3 if the box encloses the entire body of a standing person. Thus, we intentionally extend the bounding box downward so that the aspect ratio becomes 1:3 if the height is less than three times the width. Otherwise, we use the original bounding box to compute the midpoint of the bottom edge. This way, we tried to identify the standing positions of visitors if obstacles occluded their bodies.

Fig. 6 presents perspective camera views of the four sensors together with extracted bounding boxes (in orange) and manually plotted feature points (in red) at the top, and the corresponding standing positions (as orange disks) reconstructed from the computed homographies. Bounding boxes (in light blue), which enclose security guards and explainers, were intentionally excluded since their behavioral patterns were already known. For example, see a blue box in the camera image of Sensor #3 in Fig. 6.

V. Visualizing Maps of Visitors' Interest

The camera sensor data collected in the previous step could now be used to visualize the maps of interest associated with the spatial distribution of visitors in the exhibition room. For this purpose, first the entire distribution of visitors in the exhibition room was reconstructed, and then visual analysis of spatiotemporal changes in the map of interest was conducted.

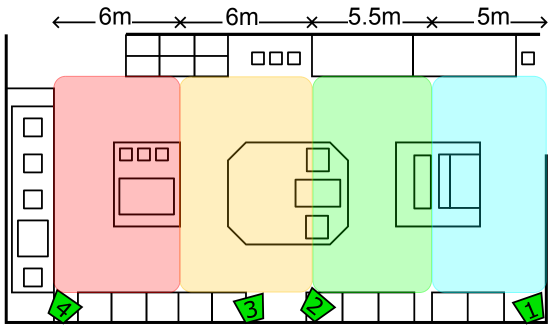

Fig. 7. Coverage area assigned to each sensor in the exhibition room.

A. Integrating Visitor Positions Obtained by Multiple Sensors

Although each of the four installed sensors was controlled to process a camera image and record the visitors' position every minute, these needed to be integrated to compose the entire distribution of the visitors in the exhibition room. Unfortunately, it was possible that two or more sensors captured the same visitors, meaning it was necessary to identify the standing position of each visitor exclusively. As shown in Fig. 5, the approach taken was to limit the coverage area of each sensor by respecting the projected area of the camera image with the corresponding homography, Fig. 7 shows the assignment of the entire exhibition room to the respective sensors.

It was also possible that some sensors failed to record the bounding boxes of visitors every minute on time or stored them multiple times within a single minute. In this case, the analysis was limited to the frames in which all four sensors successfully sent the analysis results on time and skipped other unsynchronized records. This strategy allowed synchronization of the analysis results provided by each sensor and make the analysis consistent in the spatiotemporal behaviors of visitors.

B. Visualizing the Spatiotemporal Behaviors of Visitors

After having successfully integrated the local distributions of visitors provided by multiple sensors, the entire map of interest in the exhibition room was visualized. For this purpose, the spatial distributions of visitors were transformed into visual maps respecting their temporal changes. The common idea is to trace the travel trajectories of visitors individually. However, it is not always possible to fully identify the correspondence between the visitors between adjacent temporal frames because many visitors rapidly move around in the exhibition room within one minute. Another reason is that some sets of visitors may stay still around some spots for a long time, and thus the trajectories may overlap multiple times on the floor map.

In this study, *heatmaps* were employed as the visualization tool for illustrating the spatiotemporal changes in the distribution of visitors' interest. This is beneficial because it is possible to instantly identify meaningful *hotspots* in the exhibition room through visual inspection. Furthermore, temporal changes can be incorporated into the distribution by respecting the history of visitors' positions.

The adopted approach employed the analogy of traveling visitors in the exhibition space as eye-gaze movements on the screen [26]. This was accomplished by recording the spatiotemporal behaviors of visitors in the exhibition room and computing the associated smooth density using *kernel density estimation* [32]. More specifically, the visitors' positions were convolved with Gaussian kernels as:

$$G(x, y) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{x^2 + y^2}{2}),$$

where (x, y) corresponds to 2D ground positions in the exhibition room. Suppose that (x_i, y_i) indicates the 2D ground position of the i-th visitor, where $i = 1, \ldots, L$. The heatmap is given by

$$H(x, y) = \frac{2}{L} \sum_{i=1}^{L} \frac{\max\{0, T - (t_p - t_i)\}}{T} G(x - x_i, y - y_i),$$

where t_i is the time when the visitor position was recorded, and t_p is the present time. Note that T indicates the predefined duration of the visitors' interest on which we want to focus. This helps us represent the history of their behaviors in the heatmap representation. Here, we set the predefined period to be five minutes by default. Fig. 1 shows a snapshot of the time-varying heatmap over the exhibition room, where the color changes from blue to green to red as the degree of density increases.

VI. Results

In this section, we present several experimental results to demonstrate the capability of our approach.

A. Data Acquisition from Sensors and Implementation

We recorded the measurement data obtained through the four sensors installed in the Fukushima Museum from July 2021 to February 2022. The data include bounding boxes of visitors in the camera views, which are extracted by the YOLO algorithm on the four Raspberry Pi devices every minute during the opening period of the museum. As a preprocess, we manually plotted four or more pairs of corresponding points in the camera view and floor map for each sensor to calculate the homography between them. The current approach was installed offline on workstations to check its feasibility. This program software is expected to be installed on Raspberry Pi devices later for online monitoring of the exhibition room.

B. Experimental Results

Fig. 8 shows heatmaps representing temporal sequences of maps of interest derived from the distributions of visitors in the exhibition room. Fig. 8(a) shows the spatiotemporal behaviors of individual visitors as four snapshots, in which they were likely to be interested in specific sets of exhibits in the room. Note that the snapshots correspond to the distribution of visitors every minute. The visitor in the top right of the room remained stationary for four minutes. On the other hand, visitors in the bottom left kept their attention on the collections for two minutes and left. The last snapshot reflects their past behaviors by accentuating the color of the heatmap around it. On the contrary, Fig. 8(b) presents the spatiotemporal changes in the map of interest for group visitors. The visitors were students on the educational school excursion and might have had limitations in time. Their behaviors were relatively dynamic because they often formed a small group and kept moving on to the next exhibit.

We also explored sites of interest to evaluate collection layouts in the exhibition room. For this purpose, we first summed up the sequence of heatmaps to find an accumulated map of interest for a specific period. We then normalized the grid pixel values in the accumulated map by the maximum pixel value. This facilitated us to explore meaningful hotspots in the exhibition space. Each result at the top of Fig. 9 shows an accumulated map of interest for a single day.

We also wanted to identify particular hotspots where visitors are likely to stay in the same location in the museum space for extended periods (i.e., several minutes). This inspired us to compute pixel-wise multiplications of temporally consecutive heatmaps as intermediate distribution maps and then accumulate them. The bottom row of Fig. 9 shows such maps of interest, representing sites of interest at which visitors remained for some time. Note that in this experiment, we computed the pixel-wise multiplication of six temporally consecutive heatmaps, which implies that the resulting maps of interest clarify the spatial positions where visitors were likely to stay for five minutes or more.

Fig. 9(a) indicates hotspots in the exhibition room on the free open day when the museum accommodated many groups of visitors. The top figure presents the typical distribution of interest when we integrate heatmaps for a relatively long period, for example, a month. Conversely, fewer visitors yielded smaller hotspot areas that were sparsely distributed in the exhibition space, as depicted at the top of Fig. 9(b). Fig. 9(c) at the top shows that another small set of visitors left relatively many hotspots in the map of interest, while the hotspot areas are all small and isolated. The last two cases imply that most visitors were individuals and likely to stay around the exhibits they preferred. The corresponding figures in the bottom row show the hotspot areas at which visitors were likely to stay for five or more minutes. To our surprise, the hotspots for the five-minute stay do not necessarily coincide with those we found by simply integrating all the heatmaps for the period we saw in the top figures. This helps us to discriminate sites of interest through which many visitors just passed from those where they stood still to pay particular attention.

C. Analysis and Discussions

Fig. 10 details the collections exhibited in the room associated with the Ancient (i.e., Tumulus) period. According to the hotspots we extracted from the maps of interest in the bottom row of Fig. 9, we can claim that ancient bowls and dishes made of stone, old farm tools, and items of the local Buddhist culture significantly attracted visitors to the museum.

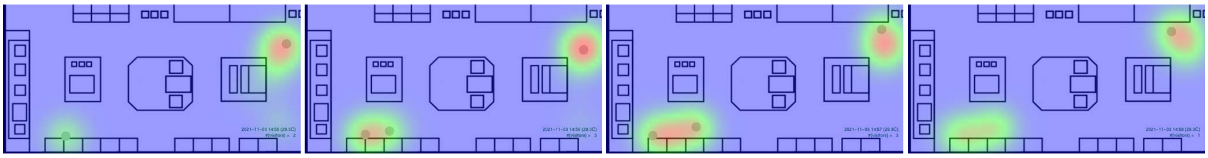

(a)

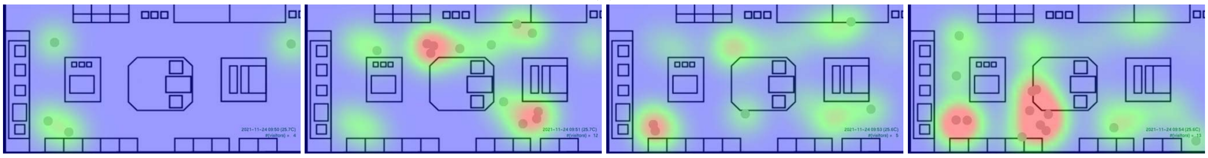

(b)

Fig. 8. Maps of interest visualized as heatmaps. Spatiotemporal changes in the map for (a) individual and (b) group visitors. The positions of visitors are marked as disks.

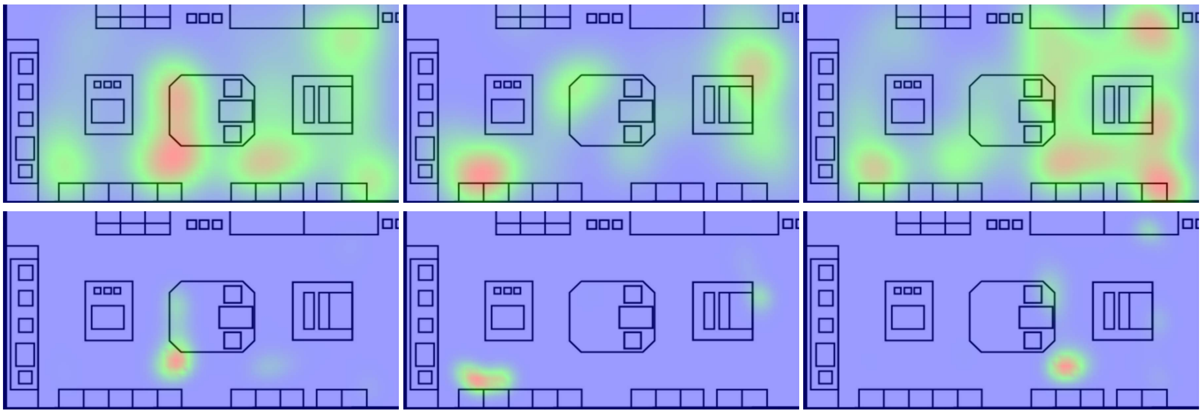

(a) August 21, 2021 (b) December 21, 2021 (c) December 23, 2021

Fig. 9. Integrals of the map of interest. Top: Accumulating maps of interest for one day. Bottom: Accumulating pixel-wise multiplication of six temporally consecutive maps for one day. (a) The map of interest presents a relatively wide distribution of visitors if the number is high. (b)(c) Locally limited distributions appeared when the museum accommodated fewer visitors.

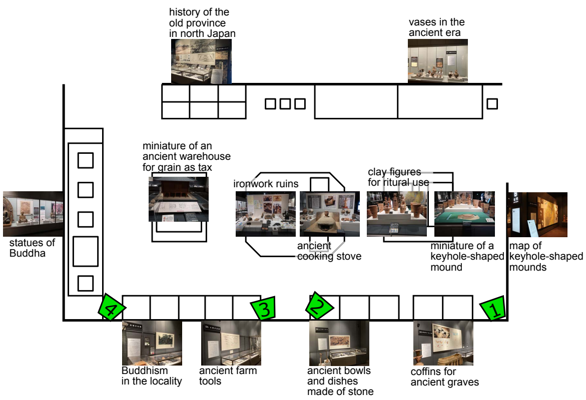

Fig. 10. Contents exhibited in the room featuring the Tumulus period.

The museum set the presentation boards in front of these three hotspots in the exhibition; hence, visitors may stand still to read the descriptions. On the other hand, the top row of Fig. 9 reveals that the area on the right of the room was relatively congested with visitors. However, they left this area within a short period, probably because this point is the starting point of the exhibition in this room, and just stopped to figure out the overall contents of the room. Another observation suggests that visitors were more interested in the exhibition of key-hole-shaped mounds (i.e., ancient Kofun graves) and ironwork ruins than other collections.

Our approach helps curators to improve the layout of collections in art galleries and museums by inferring the visitors' preferences through visual analysis. However, the proposed composition for maps of interest is still limited to visualizing the spatiotemporal positions of visitors and cannot explicitly identify which collections they intentionally view around them. Accuracy in identifying spatial positions needs to be further improved by simultaneously employing other types of sensors. Installing more sensors and devising their spatial configuration may alleviate the accuracy problem by reducing unwanted occluded areas. Implementing this approach on the Raspberry Pi computers remains to be tackled so that the curators can observe the congestion of visitors in real time.

VII. Conclusion

This paper has presented an approach to visualizing the spatiotemporal maps of interest reflecting the viewing behaviors of visitors in the exhibition space. Our technical contribution lies in the new approach for visualizing the spatial layout of visitors in the exhibition room through single-board computers equipped with sensors. This was accomplished by employing machine-learning-based object detection and homography-aware 3D reconstruction techniques. We also transformed changes in the spatiotemporal layouts of visitors as a continuous distribution map of viewing interest and visualized them as dynamic heatmaps. Visual analysis of dynamic maps of interest and their integrals over eight months of measurement data effectively clarified the positions of underlying hotspots and their associated collections the visitors most preferred.

Acknowledgments

We would like to extend our sincere gratitude to Fukushima Prefectural Museum for their invaluable assistance in providing us with the photographs and experimental environments for our research. This work has been partially supported by JSPS KAKENHI Grant Number JP19H04120.

References

[1] Y. K. Choi, "The morphology of exploration and encounter in museum layouts," *Environment and Planning B: Planning and Design*, vol. 26, no. 2, pp. 241–250, 1999. doi: 10.1068/b4525

[2] B. Hillier and K. Tzortzi, *Space Syntax*. John Wiley & Sons, Ltd, 2006, ch. 17, pp. 282–301.

[3] J. Peponis, R. C. Dalton, J. Wineman, and N. Dalton, "Measuring the effects of layout upon visitors' spatial behaviors in open plan exhibition settings," *Environment and Planning B: Planning and Design*, vol. 31, no. 3, pp. 453–473, 2004. doi: 10.1068/b3041

[4] S. S.Yalowitz and K. Bronnenkant, "Timing and tracking: Unlocking visitor behavior," *Visitor Studies*, vol. 12, no. 1, pp. 47–064, 2009. doi: 10.1080/10645570902769134

[5] S. Bitgood, "An analysis of visitor circulation: Movement patterns and the general value principle," *Curator: The Museum Journal*, vol. 49, no. 4, pp. 463–475, 2006. doi: 10.1111/j.2151-6952.2006.tb00237.x

[6] V. Kirchberg and M. Tröndle, "Experiencing exhibitions: A review of studies on visitor experiences in museums," *Curator: The Museum Journal*, vol. 55, no. 4, pp. 435–452, 2012. doi: 10.1111/j.2151-6952.2012.00167.x

[7] Y. Yoshimura, S. Sobolevsky, C. Ratti, F. Girardin, J. P. Carrascal, J. Blat, and R. Sinatra, "An analysis of visitors' behavior in the Louvre museum: A study using Bluetooth data," *Environment and Planning B: Planning and Design*, vol. 41, no. 6, pp. 1113–1131, 2014. doi: 10.1068/b130047p

[8] M. Grafe, R. Wortmann, and H. Westphal, "AR-based interactive exploration of a museum exhibit," in *Proceedings of the 1st IEEE International Workshop on Agumented Reality Toolkit*, 2002. doi: 10.1109/ART.2002.1106945

[9] F. Windhager, P. Federico, G. Schreder, K. Glinka, M. Dörk, S. Miksch, and E. Mayr, "Visualization of cultural heritage collection data: State of the art and future challenges," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 6, pp. 2311–2330, 2019. doi: 10.1109/TVCG.2018.2830759

[10] K. Dasu, K.-L. Ma, J. Ma, and J. Frazier, "Sea of genes: A reflection on visualising metagenomic data for museums," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 935–945, 2021. doi: 10.1109/TVCG.2020.3030412

[11] S. Dodge, R. Weibel, and A.-K. Lautenschütz, "Towards a taxonomy of movement patterns," *Information Visualization*, vol. 7, no. 3-4, pp. 240–252, 2008. doi: 10.1057/PALGRAVE.IVS.9500182

[12] G. Andrienko, N. Andrienko, P. Bak, D. Keim, S. Kisilevich, and S. Wrobel, "A conceptual framework and taxonomy of techniques for analyzing movement," *Journal of Visual Languages and Computing*, vol. 22, no. 3, pp. 213–232, 2011. doi: 10.1016/j.jvlc.2011.02.003

[13] N. Andrienko and G. Andrienko, "Visual analytics of movement: An overview of methods, tools and procedures," *Information Visualization*, vol. 12, no. 1, pp. 3–24, 2013. doi: 10.1177/1473871612457601

[14] Y. Gu, A. Lo, and I. Niemegeers, "A survey of indoor positioning systems for wireless personal networks," *IEEE Communications Surveys & Tutorials*, vol. 11, no. 1, pp. 13–32, 2009. doi: 10.1109/SURV.2009.090103

[15] J. Lanir, T. Kuflik, J. Sheidin, N. Yavin, K. Leiderman, and M. Segal, "Visualizing museum visitors' behavior: Where do they go and what do they do there?" *Personal and Ubiquitous Computing*, vol. 21, pp. 313–326, 2017. doi: 10.1007/s00779-016-0994-9

[16] T. Kuflik, J. Lanir, E. Dim, A. Wecker, M. Corra', M. Zancanaro, and O. Stock, "Indoor positioning: Challenges and solutions for indoor cultural heritage sites," in *Proceedings of the 16th International Conference on Intelligent User Interfaces*, ser. IUI '11, 2011, pp. 375—-378. doi: 10.1145/1943403.1943469

[17] G. Andrienko, N. Andrienko, S. Rinzivillo, M. Nanni, D. Pedreschi, and F. Giannotti, "Interactive visual clustering of large collections of trajectories," in *2009 IEEE Symposium on Visual Analytics Science and Technology*, 2009, pp. 3–10. doi: 10.1109/VAST.2009.5332584

[18] P. Bak, H. Janetzko, F. Mansmann, and D. Keim, "Spatiotemporal analysis of sensor logs using growth ring maps," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, no. 6, pp. 913–920, 2009. doi: 10.1109/TVCG.2009.182

[19] L. Chittaro, R. Ranon, and L. Ieronutti, "VU-flow: A visualization tool for analyzing navigation in virtual environments," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 6, pp. 1475–1485, 2006. doi: 10.1109/TVCG.2006.109

[20] M. Zancanaro, T. Kuflik, Z. Boger, D. Goren-Bar, and D. Goldwasser, "Analyzing museum visitors' behavior patterns," in *User Modeling 2007*, C. Conati, K. McCoy, and G. Paliouras, Eds., 2007, pp. 238–246. doi: 10.1007/978-3-540-73078-1_27

[21] E. Dim and T. Kuflik, "Automatic detection of social behavior of museum visitor pairs," *ACM Trans. Interact. Intell. Syst.*, vol. 4, no. 4, 2014. doi: 10.1145/2662869

[22] B. Breve, S. Cirillo, M. Cuofano, and D. Desiato, "Enhancing spatial perception through sound: Mapping human movements into midi," *Multimedia Tools and Applications*, vol. 81, pp. 73–94, 2022. doi: 10.1007/s11042-021-11077-7

[23] G. Andrienko, N. Andrienko, M. Burch, and D. Weiskopf, "Visual analytics methodology for eye movement studies," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 12, pp. 2889–2898, 2012. doi: 10.1109/TVCG.2012.276

[24] M. Raschke, T. Blascheck, and M. Burch, *Visual Analysis of Eye Tracking Data*. Springer New York, 2014, pp. 391–409.

[25] T. Blascheck, K. Kurzhals, M. Raschke, M. Burch, D. Weiskopf, and T. Ertl, "State-of-the-art of visualization for eye tracking data," in *EuroVis - STARs*, R. Borgo, R. Maciejewski, and I. Viola, Eds. The Eurographics Association, 2014, pp. 63–82. doi: 10.2312/eurovisstar.20141173

[26] S. Takahashi, A. Uchita, K. Watanabe, and M. Arikawa, "Gaze-driven placement of items for proactive visual exploration," *Journal of Visualization*, vol. 25, pp. 613–633, 2022. doi: 10.1007/s12650-021-00808-5

[27] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788. doi: 10.1109/CVPR.2016.91

[28] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," 2018, doi: 10.48550/arXiv.1804.02767.

[29] D. Liebowitz, A. Criminisi, and A. Zisserman, "Creating architectural models from images," *Computer Graphics Forum*, vol. 18, no. 3, pp. 39–50, 1999. doi: 10.1111/1467-8659.00326

[30] Z. Zhang, "Flexible camera calibration by viewing a plane from unknown orientations," in *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 1, 1999, pp. 666–673. doi: 10.1109/ICCV.1999.791289

[31] A. Agarwal, C. Jawahar, and P. Narayanan, "A survey of planar homography estimation techniques," International Institute of Information Technology Hyderabad, India, Tech. Rep. Centre for Visual Information Technology, Tech. Rep. IIIT/TR/2005/12, 2005.

[32] O. Daae Lampe and H. Hauser, "Interactive visualization of streaming data with kernel density estimation," in *2011 IEEE Pacific Visualization Symposium*, 2011, pp. 171–178. doi: 10.1109/PACIFICVIS.2011.5742387